

REMARKS

By this amendment, no claims have been amended, cancelled, or added. Hence, Claims 1-23 and 49-71 are pending in the application.

THE CLAIM OBJECTIONS/REJECTIONS

Claims 2-5, 12, 50-53, and 60 are objected to as being dependent upon a rejected base claim, but would be allowable if rewritten in independent form.

Claims 1, 10, 18, 20-23, 49, 58, 66, and 68-71 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over U.S. Patent No. 6,453,404 issued to Bereznyi et al. ("*Bereznyi*") in view of U.S. Patent Number 6,453,356 issued to Sheard et al. ("*Sheard*"). Claims 6-9 and 54-57 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over *Bereznyi* in view of *Sheard* in view of U.S. Patent Number 5,659,682 issued to Devarakonda et al. ("*Devarakonda*"). Claims 13 and 61 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over *Bereznyi* in view of *Sheard* in view of *Devarakonda* in view of U.S. Patent Number 6,032,188 issued to Maris et al. ("*Maris*") in view of U.S. Patent Number 6,243,814 issued to Matena et al. ("*Matena*"). Claims 14, 16, 62, and 64 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over *Bereznyi* in view of *Sheard* in view of *Devarakonda* in view of *Maris*. Claims 15 and 63 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over *Bereznyi* in view of *Sheard* in view of *Devarakonda* in view of *Maris* in view of U.S. Patent Number 6,012,085 issued to Yohe et al. ("*Yohe*"). Claims 17 and 65 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over *Bereznyi* in view of *Sheard* in view of *Devarakonda* in view of *Maris* in view of U.S. Patent Number 5,933,849 issued to Srblic et al. ("*Srblic*"). Claims 11 and 59 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over *Bereznyi* in view of *Sheard* in view of U.S. Patent Number 5,999,940 issued to Ranger et al. ("*Ranger*").

Claims 19 and 67 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over *Bereznyi* in view of *Sheard* in view of U.S. Patent Number 6,832,120 issued to Frank et al. (“*Frank*”).

Applicants respectfully traverse.

CLAIMS 1-23 AND 49-71 ARE PATENTABLE OVER THE CITED ART

As explained in detail below, even if the cited art could be properly combined, the resulting combination would not disclose, teach, or suggest at least one element in each pending claim, due to the significant, fundamental differences between the approach of the pending claims and the approach of the nine cited art references, and in particular, the approaches of *Bereznyi* and *Sheard*.

The features of Claim 1

Claim 1 recites the following features:

“A method for managing caches in a system with multiple caches that may contain different copies of a data item, comprising the steps of:
modifying the data item in a first node of said multiple caches to create a modified data item;
sending the modified data item from said first node to a second node of said multiple caches without durably storing the modified data item from said first node to persistent storage;
after said modified data item has been sent from said first node to said second node, said first node sending a request to a master of said data item for writing said data item to persistent storage; and
in response to said request, said master coordinating with said multiple caches to cause said data item to be written to persistent storage”

The above-combination of elements is not disclosed, taught, or suggested, either individually or in combination, by *Bereznyi* or *Sheard*.

The Applicants’ specification describes various techniques for coordinating the writing of dirty data items to persistent storage in systems that allow a dirty version of the same data item to reside in multiple caches. As taught in the Applicants’ specification in paragraph 8, a

dirty data item is a data item that (1) contains changes, and (2) has not yet been persistently stored.

Claim 1 is directed towards a method for managing caches in a system with multiple caches. The system of multiple caches may contain different copies of a data item in the multiple caches. Initially, a data item, in a first node of the multiple caches, is modified to create a modified data item. Next, the modified data item is sent from the first node to a second node of the multiple caches without durably storing the modified data item from the first node to persistent storage.

After the modified data item has been sent from the first node to the second node, the first node sends a request, to a master of the data item, for writing the data item to persistent storage. In response to the request, the master coordinates with the multiple caches to cause the data item to be written to persistent storage. For example, the master may (1) grant the first node permission to perform the write, or (2) inform the first node that another node has already written to persistent storage a version that is at least as recent as the dirty version stored in the requesting node.

Advantageously, the approach of Claim 1 allows dirty data items to be transferred between caches without first writing the dirty data item to persistent storage. In other words, the approach of Claim 1 allows a data item, which (1) contains changes, and (2) has not yet been persistently stored, to be transferred between caches without first writing the data item to persistent storage.

The teachings of the cited art

While the approach of the pending claims is directed towards transferring a modified data item between caches without first writing the modified data item to persistent storage, the

approaches of the cited art references are directed towards sharply different subject matter. As a result, fundamental differences exist between the subject matter of the cited art references and the approach of the pending claims.

For example, *Bereznyi* describes an approach for allocating memory for storage of data items by defining a series of small blocks of a cache that are uniform in size. Blocks of memory are allocated based on the size of the data item to be stored in the cache (Abstract). FIG. 5 of *Bereznyi*, and the corresponding description, also describe a multi-level cache wherein an attempt to retrieve a requested data item is initially made from a local cache 120 (the L1 cache). If the local cache 120 does not have the requested data item, then an attempt is made to retrieve the requested data item from the remote cache 126 (L2 cache). Finally, if the remote cache 126 does not have the requested data item, then an attempt is made to retrieve the requested data item from the content provider 136 (L3 cache). Once a requested data item is retrieved from a higher level cache, then the retrieved data item may be stored in a lower level cache.

In the approach of *Bereznyi*, the client 102, the local cache 120, and the remote caches 126 do not modify any data items. Instead, the purpose of the multi-level cache of FIG. 5 is to provide unmodified data items to the client 102 faster than if the unmodified data items were provided from the content provider 136. Thus, while data items may be transferred from one cache to another in the approach of *Bereznyi*, *Bereznyi* lacks any teaching of transferring a modified data item that has not yet been persistently stored.

Importantly, there are numerous fundamental differences between the approach of *Bereznyi* and that of Claim 1. For example, *Bereznyi* does not contain any suggestion of sending a modified data item from one cache to another cache without durably storing the modified data item. Further, the teaching of *Bereznyi* lacks anything remotely analogous to a

master coordinating with multiple caches to cause a data item to be written to persistent storage.

The *Sheard* reference also has nothing to do with transferring a modified data item from one cache to another. Instead, *Sheard* is directed towards an approach for exchanging data between two or more applications. A technology dependent data stream, originating from a source application, is transformed into a technology independent stream using an adapter. A data exchange engine receives the technology independent stream and identifies a destination application. The technology independent stream is then transformed into a technology dependent stream, associated with the destination application, using a different adapter.

Bereznyi and Sheard, individually or in combination, fail to show numerous features of Claim 1

As stated in MPEP § 2143.03: “To establish prima facie obviousness of a claimed invention, all the claim limitations must be taught or suggested by the prior art.” In *Re Royka*, 180 USPQ 580.

Applicant admits to being perplexed about how to respond to the inconsistency between (1) the evidence required to support an obviousness rejection, and (2) the evidence that has been offered in the Office Action relating to the present application. Specifically, to support an obviousness rejection, the Applicant would expect an argument that has the following form: (1) element X is shown in reference A, (2) element Y is shown in reference B, and (3) there is some actual suggestion to combine the references A and B to create the mechanism or technique that has both elements X and Y.

However, the Office Action does not support the obviousness rejections in that manner. Rather, to support the obviousness rejections, not only has each claim been divided into its

constituent elements, but also each constituent element of the claim has been finely dissected into a set of short phrases and sentence fragments. The Office Action then points out how each individual fragment corresponds to a similar fragment in any one of a handful of references. The fragment-to-prior-art correlation appears to have been made without any consideration as to the relationship between the fragments, the meaning of the elements as a whole, and the meaning of the claim as a whole.

By way of example, consider a claim to “a device for blowing out one or more candles on a birthday cake in response to being presented the birthday cake.” By the Office Action’s logic, this claim may be rejected based on a combination of references discussing: (a) a hurricane – which shows “...blowing...,” (b) a 18th century Spanish Mission which was illuminated by a plurality of prayer candles on a candelabra – to show “...one or more candles on a...,” (c) a police report describing a car accident – to show “...in response to...” because the accident happened in response to two cars colliding, (d) the Wheel of Fortune® game show – because letters are “...presented...” to game contestants, (e) a recipe for a cherry sauce to accompany an Angel Food cake – which shows “...cake,” (f) an obituary of Grover Cleveland, the 22nd and 24th President of the United States of America – which includes a mention of the day of his birth to show “...birthday....,” and (g) a copy of a Popular Mechanics magazine from March 1984, which discusses “a device for...” It is respectfully submitted that these references, either individually or in combination, do not render obvious “a device for blowing out one or more candles on a birthday cake in response to being presented with the birthday cake.”

In the same way, the nine cited references have been cited to show fragments of the claim elements, without any regard to the meaning ascribed to those claim fragments given their placement in the claims.

The Office Action acknowledges that *Bereznyi* “does not teach the sending of requests, the responding to requests, and the writing of data to persistent storage.” Whether or not this is true, the Office Action relies upon various portions of *Bereznyi* to show fragments of various elements of Claim 1. For example, *Bereznyi* is cited to show the fragments of (1) “...modifying the data item in a first node of said multiple caches...” at Col. 21, lines 22-25; Col. 4, lines 55-60; and Col. 6, lines 53-55, (2) “...to create a modified data item...” at Col. 5, lines 43-46 and Col. 21, lines 2-25, (3) “...sending the modified data item from said first node to a second node of said multiple caches...” at Col. 23, lines 15-16; Col. 21, lines 22-25; Col. 4, lines 55-60; and Col. 6, lines 53-55, (4) “...without durably storing the modified data item from said first node to persistent storage...” at Col. 1, lines 16-18; Col. 21, lines 22-25; Col. 4, lines 55-60, (5) “...after said modified data item has been sent from said first node to said second node...” at Col. 21, lines 22-25; Col. 23, lines 15-16; and Col. 4, lines 55-60, (6) “...to a master of said data item...” at Col. 8, lines 27-29 and Col. 21, lines 22-25, and (7) “...said master coordinating with said multiple caches...” at Col. 8, lines 27-29.

However, no portion of *Bereznyi* shows, or is cited to show, the elements in Claim 1 of:

“sending the modified data item from said first node to a second node of said multiple caches without durably storing the modified data item from said first node to persistent storage;” or

“after said modified data item has been sent from said first node to said second node, said first node sending a request to a master of said data item for writing said data item to persistent storage;” or

“in response to said request, said master coordinating with said multiple caches to cause said data item to be written to persistent storage”

Thus, even though the cited art references may use certain words or phrases that are also in Claim 1, the meaning ascribed to those words or phrases, given their placement in Claim 1,

is not disclosed, taught, or suggested by the cited art. Further, as explained below, the cited art references often fail to show the words or phrases they are alleged to show.

To illustrate, the portion of *Bereznyi* relied upon by the Office Action to show the claim fragment of "...modifying the data items in a first node of said multiple caches..." states, *in toto*:

For example, the user may get a particular data item's properties from the cache and modify them to alter the time and/or frequency limit data and return the edited properties to the cache. **Thus, the data item is unaltered**, but its properties have been altered... The remote computer 40 may be another personal computer, a server, a router, a network PC, a peer device or other common network node, and typically includes many or all of the elements described above relative to the personal computer 20, although only a memory storage device 42 has been illustrated in FIG. 1... With the use of multiple caches, the system 100 advantageously permits distributed caching, as illustrated in FIG. 5. (Col. 21, lines 22-25, Col. 4, lines 55-60, and Col. 6, lines 53-55). (emphasis added).

This cited portion of *Bereznyi*, instead of discussing modifying the data items in a first node of multiple caches, merely discusses (a) modifying the properties of a data item, while leaving the data item itself unaltered, e.g., the properties of a data item may be used to determine when to remove the data item from the cache, (b) remote computer 40 may be a variety of devices, and (c) FIG. 5 shows a multi-level cache. These topics have nothing to do with modifying data items within a cache. Importantly the portion cited by the Office Action clearly states, "the data item is unaltered." Consequently, the portion cited by the Office Action expressly teaches away from the Office Action's suggestion that the *Bereznyi* teaches the claim fragment of "modifying the data items in a first node of said multiple caches."

Further, the Office Action alleges that *Bereznyi* teaches the use of a master as claimed in Claim 1. However, the portion of *Bereznyi* cited to show a master, featured in the element of "in response to said request, said master coordinating with said multiple caches to cause said data item to be written to persistent storage," merely states, *in toto*:

The main cache class is CCache, which is responsible for managing the local cache, one or more remote cache servers, and all data that flows between them. (Col. 8, lines 27-29).

The above portion of *Bereznyi* lacks any suggestion of a master, in response to a request for writing the data item to persistent storage, which causes a data item to be written to persistent storage. While the cited portion discusses that the main cache class is responsible for managing different caches, this act is not analogous to “in response to said request, said master coordinating with said multiple caches to cause said data item to be written to persistent storage” because, *inter alia*, the cited portion of *Bereznyi* lacks any mention of persistently storing data items.

Thus, there are significant fundamental differences between what *Bereznyi* actually teaches and what it is relied upon to show by the Office Action. Since *Bereznyi* does not teach (a) sending a modified data item from a first node to a second node of multiple caches without durably storing the modified data item, (b) a first node sending a request to a master of the data item for writing the data item to persistent storage, and (c) a master coordinating with multiple caches to cause the data item to be written to persistent storage, *Bereznyi* cannot possibly disclose, teach, or suggest the following elements of Claim 1:

“sending the modified data item from said first node to a second node of said multiple caches without durably storing the modified data item from said first node to persistent storage;” or

“after said modified data item has been sent from said first node to said second node, said first node sending a request to a master of said data item for writing said data item to persistent storage;” or

“in response to said request, said master coordinating with said multiple caches to cause said data item to be written to persistent storage”

Similarly, *Sheard* is cited to show numerous other claim fragments of Claim 1, including “persistent storage,” “said first node sending a request,” “for writing said data item to

persistent storage,” “and in response to said request,” and “to cause said data item to be written to persistent storage.” The Applicants readily agree that (a) forms of persistent storage existed prior to their invention, and (b) requests were sent from nodes prior to their invention.

However, Applicants respectfully submit that *Sheard* is directed towards sharply different subject matter than that featured in Claim 1. For example, at a high level, *Sheard* lacks any suggestion of sending a modified data item from a first cache to a second cache. Consequently, the teachings of *Sheard* do not cure the deficiencies of *Bereznyi* in teaching or suggesting the elements of Claim 1.

Thus, neither *Bereznyi* nor *Sheard* individually teach each the following claim limitations:

“sending the modified data item from said first node to a second node of said multiple caches without durably storing the modified data item from said first node to persistent storage;
after said modified data item has been sent from said first node to said second node, said first node sending a request to a master of said data item for writing said data item to persistent storage; and
in response to said request, said master coordinating with said multiple caches to cause said data item to be written to persistent storage”

As a result of the numerous fundamental differences between (a) what *Bereznyi* teaches, (b) what *Sheard* teaches, and (c) what is featured in Claim 1, even if *Bereznyi* and *Sheard* were to be properly combined, the resulting combination would still not disclose, teach, or suggest numerous elements of Claim 1. Consequently, it is respectfully submitted that Claim 1 is patentable over the cited art, and is in condition for allowance.

Claims 2-23 and 49-71 are patentable over the cited art

Independent Claim 49 recites features similar to that of Claim 1, except that Claim 49 is recited in computer-readable medium format. Consequently, it is respectfully submitted that,

for at least the reasons given above with respect to Claim 1, Claim 49 is patentable over the cited art and is in condition for allowance.

Claims 2-23 and 50-71 are dependent claims, each of which depends (directly or indirectly) on one of the claims discussed above. Each of Claims 2-23 and 50-71 are therefore allowable for the reasons given above for the claim on which it depends. In addition, each of Claims 2-23 and 50-71 introduces one or more additional limitations that independently render it patentable. However, due to the fundamental differences already identified, to expedite the positive resolution of this case a separate discussion of those limitations is not included at this time, although the Applicants reserve the right to further point out the differences between the cited art and the novel features recited in the dependent claims.

CONCLUSION

For the reasons set forth above, it is respectfully submitted that all of the pending claims are now in condition for allowance. Therefore, the issuance of a formal Notice of Allowance is believed next in order, and that action is most earnestly solicited.

The Examiner is respectfully requested to contact the undersigned by telephone if it is believed that such contact would further the examination of the present application.

To the extent necessary, a petition for an extension of time under 37 C.F.R. § 1.136 is hereby made. Please charge any fee shortages or credit any overages to Deposit Account No. 50-1302.

Respectfully submitted,
HICKMAN PALERMO TRUONG & BECKER LLP



Christopher J. Brokaw
Reg. No. 45,620

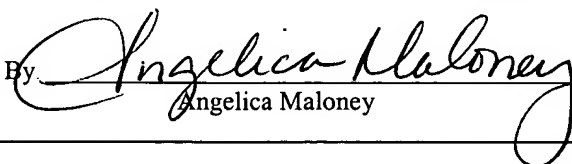
2055 Gateway Place, Suite 550
San Jose, CA 95110-1089
(408) 414-1080 ext. 225
Date: **October 7, 2005**
Facsimile: (408) 414-1076

CERTIFICATE OF MAILING

I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail in an envelope addressed to: **Mail Stop Amendment**, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

On October 7, 2005

By


Angelica Maloney